arXiv:1401.6227v1 [cs.PL] 24 Jan 2014

From Safety To Termination And Back: SMT-Based Verification For Lazy Languages

Niki Vazou Eric L. Seidel Ranjit Jhala
UC San Diego

Abstract

SMT-based verifiers have long been an effective means of ensuring safety properties of programs. While these techniques are well understood, we show that they implicitly require *eager* semantics; directly applying them to a *lazy* language is unsound due to the presence of divergent sub-computations. We recover soundness by composing the safety analysis with a *termination* analysis. Of course, termination is itself a challenging problem, but we show how the safety analysis can be used to ensure termination, thereby bootstrapping soundness for the entire system. Thus, while safety invariants have long been required to prove termination, we show how termination proofs can be to soundly establish safety. We have implemented our approach in LIQUIDHASKELL, a Refinement Type-based verifier for Haskell. We demonstrate its effectiveness via an experimental evaluation using LIQUIDHASKELL to verify safety, functional correctness and termination properties of real-world Haskell libraries, totaling over 10,000 lines of code.

1. Introduction

SMT-based verifiers, based on Floyd-Hoare Logic (*e.g.* EscJava [12]), or combined with abstract interpretation (*e.g.* SLAM [1]), have been highly effective at the automated verification of imperative programs. In the functional setting, these techniques are generalized as *refinement types*, where invariants are encoded by composing types with SMT-decidable refinement predicates [28, 38]. For example

```
type Pos = {v:Int | v >  0}
type Nat = {v:Int | v >= 0}
```

are the basic type `Int` refined with logical predicates that state that "the values" `v` described by the type are respectively strictly positive and non-negative. We encode *pre-* and *post*-conditions (contracts) using refined function types like

```
div :: n:Nat -> d:Pos -> {v:Nat | v <= n}
```

which states that the function `div` *requires* inputs that are respectively non-negative and positive, and *ensures* that the output is less than the first input `n`. If a program containing `div` statically typechecks, we can rest assured that executing the program will not lead to any unpleasant divide-by-zero errors. Several groups have demonstrated that refinements can be used to statically verify properties ranging from simple array safety [27, 38] to functional correctness of data structures [18], security protocols [3, 5], and compiler correctness [13].

Given the remarkable effectiveness of the technique, we embarked on the project of developing a refinement type based verifier for Haskell, assuming that the standard soundness proofs from Floyd-Hoare logics and refinement types would carry over directly. Of course, the previous logics and systems were all developed for eager, *call-by-value* languages, but we presumed that the order of evaluation would surely prove irrelevant, and that the soundness guarantees would translate to Haskell's lazy, *call-by-need* regime.

To our surprise, we were totally wrong.

1. Laziness Precludes Partial Correctness Our first contribution is to demonstrate that refinement typing is *unsound* in the presence of lazy evaluation (§ 2). Consider the program:

```
foo     :: n:Nat -> {v:Nat | v < n}
foo n   = if (n > 0) then (n-1) else (foo n)

bar     :: z:Pos -> x:Int -> Int
bar z x = 2014 `div` z

main    = let (a,b) = (0, foo 0) in bar a b
```

A standard refinement type checker will happily verify the above program. The refinement type signature for `foo` captures the *partial correctness* property: the function `foo` requires non-negative inputs and ensures that its output (*if* one is produced!) will be strictly less than its input. Consequently, the checker concludes that at the call-site for `bar`, the value `z` equals 0 and the value `x` is some non-negative integer that is strictly less than 0. In other words, the checker concludes that the environment is *inconsistent* and hence trivially type checks.

The issue is independent of refinement typing and affects any Floyd-Hoare logic-based verifier. For example, a hypothetical ESC-Scala would verify the following code which restates the above using classical `requires` and `ensures` clauses:

```
def foo(n:Int):Int =
  //@ requires (0 <= n)
  //@ ensures  (0 <= \result && \result < n)
  if (n > 0) return (n-1) else return foo(n)

def bar(z:Int, x:=>Int) =
  //@ requires 0 < z
  return (2014 div z)

def main = {val (a,b) = (0,foo(0));bar(a,b)}
```

One should not be alarmed as this deduction is perfectly sound under eager, *call-by-value* semantics. In both cases, the verifier determines that the call to `bar` is *dead code* – the call is safe because it is not invoked at all. This reasoning is quite unsound for Haskell's lazy, *call-by-need* semantics, and Scala's lazy, *call-by-name* parameters (indicated by the `:=>` type annotation). In both cases, the program execution would skip blithely over the call to `foo`, plunge headlong into the `div`, and crash.

As we show, the problem is that with lazy evaluation, one can only trust a refinement or invariant if one is guaranteed that evaluating the corresponding term will *not diverge*. That is, the classical

Floyd-Hoare separation of "partial" and "total" correctness breaks down, and even safety verification requires checking termination.

2. Termination for and by Refinement Typing The prognosis seems dire: to solve one problem it appears we must first solve a harder one! Our second contribution is to demonstrate that refinement types can themselves be used to prove termination (§ 2). In particular, we show how to adapt the classical idea of ranking functions [32], as embodied via sized types [2, 15], to the setting of Refinement Typing. The key idea is to ensure that each recursive call is made with parameters of strictly decreasing `Nat`-valued size.

We show that refinements naturally encode sized types and generalize them in useful ways by allowing (1) different notions of size to account for recursive data types, (2) lexicographically ordered ranking functions to support complex forms of recursion, and most importantly (3) the use of auxiliary relational invariants (circularly, via refinement types!) to verify that sizes decrease in non-structurally recursive functions.

Our use of refinements to prove termination makes proving soundness interesting in two ways. First, to check Haskell codebases, we cannot require that *every* term terminates, and so we must support programs containing terms that may diverge. Second, there is a circularity in the soundness proof itself as termination is required to prove refinements sound and vice versa. We address these issues by developing a core calculus (§ 3) with *optimistic* semantics [11], proving those semantics equivalent to call-by-name evaluation, and then proving soundness with respect to the optimistic semantics (§ 4). Thus, while it is well known that safety properties (invariants) are needed to prove termination [8], we show for the first time how termination properties are needed to prove safety.

3. Refinement Types for Real-World Haskell We have implemented our technique by extending LIQUIDHASKELL [34]. Our third contribution is an experimental evaluation of the effectiveness of our approach by using LIQUIDHASKELL to check substantial, real-world Haskell libraries totaling over $10,000$ lines of code (§ 5). The verification of these libraries requires precisely analyzing recursive structures like lists and trees, tracking the relationships between their contents, inferring invariants in the presence of low-level pointer arithmetic, and lifting the analysis across polymorphic, higher-order functions. We demonstrate that by using LIQUIDHASKELL we were able to prove termination and a variety of critical safety and functional correctness properties with a modest number of manually specified hints, and even find and fix a subtle correctness bug related to unicode handling in TEXT. To the best of our knowledge, this is the most substantial evaluation of refinement types on third party code, and demonstrates that SMT-based safety- and termination-verification can be practical and effective for higher-order, functional languages.

2. Overview

We start with an overview of our contributions. After quickly recapitulating the basics of *refinement types* we illustrate why the approach is unsound in the presence of lazy evaluation. Next, we show that we can recover soundness by coupling refinements with a termination analysis. Fortunately, we demonstrate how we can bootstrap off of refinements to solve the termination problem, and hence obtain a sound and practical verifier for Haskell.

SMT-Based Refinement Type Checking Recall the refinement type aliases `Pos` and `Nat` and the specification for `div` from § 1. A refinement type system will use these to *reject*

```
bad :: Nat -> Nat -> Int
bad x y = x `div` y
```

because, to check that the second parameter `y` has type `Pos` at the call to `div`, the system will issue a subtyping query

$$\mathtt{x} : \{\mathtt{x} \geq 0\},\ \mathtt{y} : \{\mathtt{y} \geq 0\} \vdash \{\mathtt{v} = \mathtt{y}\} \preceq \{\mathtt{v} > 0\}$$

which reduces to the *invalid* SMT query

$$(\mathtt{x} \geq 0) \wedge (\mathtt{y} \geq 0) \Rightarrow (\mathtt{v} = \mathtt{y}) \Rightarrow (\mathtt{v} > 0)$$

On the other hand, the system will *accept* the program

```
good :: Nat -> Nat -> Int
good x y = x `div` (y + 1)
```

Here, the corresponding subtyping query is

$$\mathtt{x} : \{\mathtt{x} \geq 0\},\ \mathtt{y} : \{\mathtt{y} \geq 0\} \vdash \{\mathtt{v} = \mathtt{y} + 1\} \preceq \{\mathtt{v} > 0\} \quad (1)$$

which reduces to the *valid* SMT query

$$(\mathtt{x} \geq 0) \wedge (\mathtt{y} \geq 0) \Rightarrow (\mathtt{v} = \mathtt{y} + 1) \Rightarrow (\mathtt{v} > 0) \quad (2)$$

2.1 Laziness Makes Refinement Typing Unsound

Next, let us look in detail at the program with `foo` and `bar` from § 1. A standard refinement type checker first verifies the signature of `foo` in the classical rely-guarantee fashion, by (inductively) assuming its type and checking that its body yields the specified output. In the **`then`** branch, the output subtyping obligation

$$\mathtt{n} : \{\mathtt{n} \geq 0\},\ _ : \{\mathtt{n} > 0\} \vdash \{\mathtt{v} = \mathtt{n} - 1\} \preceq \{0 \leq \mathtt{v} \wedge \mathtt{v} < \mathtt{n}\}$$

reduces to the valid SMT formula

$$\mathtt{n} \geq 0 \ \wedge\ \mathtt{n} > 0 \Rightarrow (\mathtt{v} = \mathtt{n} - 1) \Rightarrow\ (0 \leq \mathtt{v} \wedge \mathtt{v} < \mathtt{n})$$

In the **`else`** branch, the output is proven by inductively using the assumed type for `foo`. Next, inside `main` the binder `b` is assigned the output type of `foo` with the formal `n` replaced with the actual `0`. Thus, the subtyping obligation at the call to `bar` is

$$\mathtt{a} : \{\mathtt{a} = 0\},\ \mathtt{b} : \{0 \leq \mathtt{b} \wedge \mathtt{b} < 0\} \vdash \{\mathtt{v} = \mathtt{a}\} \preceq \{\mathtt{v} > 0\}$$

which reduces to the SMT query

$$\mathtt{a} = 0 \ \wedge\ (0 \leq \mathtt{b} \wedge \mathtt{b} < 0) \Rightarrow (\mathtt{v} = \mathtt{a}) \Rightarrow (\mathtt{v} > 0)$$

which is trivially valid as the antecedent is inconsistent.

Unfortunately, this inconsistency is unsound under Haskell's lazy evaluation! Since `b` is not required, the program will dive headlong into evaluating the `div` and hence crash, rendering the type checker's guarantee meaningless.

Reconciling Laziness And Refinements One may be tempted to get around the unsoundness via several different routes. First, one may be tempted to point the finger of blame at the "inconsistency" itself. Unfortunately, this would be misguided, since such inconsistencies are not a bug but a crucial feature of refinement type systems. They enable, among other things, *path sensitivity* by incorporating information from run-time tests (guards) and hence let us verify that expressions that throw catastrophic exceptions (*e.g.* `error e`) are indeed unreachable dead code and will not explode at run-time. Second, one might use a CPS transformation [25, 36] to convert the program into call-by-value. We confess to be somewhat wary of the prospect of translating inferred types and errors *back* to the source level after such a transformation. Previous experience shows that the ability to map types and errors to source is critical for usability. Third, one may want some form of *strictness analysis* [20] to statically predict which expressions must be evaluated, and only use refinements for those expressions. This route is problematic as it is unclear whether one can develop a sufficiently precise strictness analysis. More importantly, it is often useful to add *ghost* values into the program for the sole purpose of making refinement

types *complete* [33]. By construction these values are not used by the program, and would be thrown away by a strictness analysis, thus precluding verification.

2.2 Ensuring Soundness With Termination

The crux of the problem is that when we establish that

$$\Gamma \vdash \texttt{e} : \{\texttt{v:Int} \mid \texttt{p}\}$$

what we have guaranteed is that *if* `e` reduces to an integer n, *then* n satisfies the logical predicate `p` $[n/\texttt{v}]$ [14, 19]. Thus, to account for diverging computations, we should properly view the above typing judgment as weakened with a *bottom disjunct*

$$\Gamma \vdash \texttt{e} : \{\texttt{v:Int} \mid \texttt{v} = \bot \vee \texttt{p}\}$$

Now, consider the expression `let x = e in e'`. In an eager setting, we can readily eliminate the bottom disjunct and assume that `x` satisfies `p[x/v]` when analyzing `e'` because if `e` diverges, then `e'` is *not evaluated*. In other words, the mere fact that evaluation of `e'` *began* allows us to conclude that $\texttt{x} \neq \bot$, and so we can eliminate the bottom disjunct without compromising soundness. However, in a lazy setting we cannot drop the bottom disjunct because we may well evaluate `e'` even if `e` diverges!

One way forward is to take the bull by the horns and *directly* reason about divergence and laziness using the bottom disjunct. That is, to *weaken* each refinement with the bottom disjunct. While sound, such a scheme is imprecise as the hypotheses will be too weak to let us prove interesting *relational* invariants connecting *different* program variables. For instance, the subtyping query for `good` (1), if we ignore the `x` binder, becomes:

$$\texttt{y} : \{\texttt{y} = \bot \vee \texttt{y} \geq 0\} \vdash \{\texttt{v} = \bot \vee \texttt{v} = \texttt{y} + 1\} \preceq \{\texttt{v} = \bot \vee \texttt{v} > 0\}$$

which boils down to the SMT query

$$(\texttt{y} = \bot \vee \texttt{y} \geq 0) \Rightarrow (\texttt{v} = \bot \vee \texttt{v} = \texttt{y} + 1) \Rightarrow (\texttt{v} = \bot \vee \texttt{v} > 0)$$

which is invalid, causing us to reject a perfectly safe program!

One might try to make the direct approach a little less naïve by somehow axiomatizing the semantics of operators like `+` to stipulate that the result (*e.g.* `v` above) is only non-bottom when the operands (*e.g.* `y` and `1`) are non-bottom. In essence, such an axiomatization would end up encoding lazy evaluation inside the SMT solver's logic, and would indeed make the above query valid. However, it is quite unclear to us how to design such an axiomatization in a systematic fashion. Worse, even with such an axiomatization, we would never be able to verify trival programs like

```
baz :: x:Int -> y:Int -> {z:Int | x > y}
    -> Int
baz x y z = assert (x > y) 0
```

as the bottom disjunct on `z` – which is never evaluated – would preclude the verifier from using the refinement relating the values of `x` and `y` that is needed to prove the assertion. The above example is not contrived; it illustrates a common idiom of using ghost variables that carry "proofs" about other program variables.

The logical conclusion of the above line of inquiry is that to restore soundness and preserve precision, we need a means of precisely eliminating the bottom disjunct, *i.e.* of determining when a term is definitely *not* going to diverge. That is, we need a termination analysis. With such an analysis, using Reynolds' [26] terminology, we could additionally type each term as either trivial, meaning it *must* terminate, or serious, meaning it *may not* terminate. Furnished with this information, the refinement type checker may soundly drop the bottom disjunct for trivial expressions, and keep the bottom disjunct (or just enforce the refinement `true`) for serious terms.

Our approach is properly viewed as an optimization that strengthens the "direct" refinement (with a bottom disjunct) with a termination analysis that lets us eliminate the bottom disjunct in the common case of terminating terms. In general, this approach leaves open the possibility of directly reasoning about bottom (*e.g.* when reasoning about infinite streams), as we *do not* require that *all* terms be provably terminating, but only the ones with (non-trivial) refinements *without* the bottom disjunct. The dual approach – using safety invariants to strengthen and prove termination is classical [8]; this is the first time termination has been used to strengthen and prove safety!

As an aside, readers familiar with fully dependently typed languages like Agda [21] and Coq [4] may be unsurprised at the termination requirement. However, the role that termination plays here is quite different. In those settings, arbitrary terms may appear in types; termination ensures the semantics are well-defined and facilitates type equivalence. In contrast, refinement logics are carefully designed to preclude arbitrary terms; they only allow logical predicates over well-defined, decidable theories, which crucially *also* include program variables. As we saw, this choice is sound under call-by-value, but problematic under call-by-name, as in the latter setting even a mere first-order variable can correspond to an undefined diverging computation.

2.3 Ensuring Termination With Refinements

How shall we prove termination? Fortunately, there is a great deal of research on this problem. The heart of almost all the proposed solutions is the classical notion of *ranking functions* [32]: in any potentially *looping* computation, prove that some well-founded metric strictly decreases every time around the loop.

Sized Types In the context of typed functional languages, the primary source of looping computations is recursive functions. Thus, the above principle can be formalized by associating a notion of *size* with each type, and verifying that in each recursive call, the function is invoked with arguments whose size is *strictly smaller* than the current inputs. Thus, one route to recovering soundness would be to perform a first phase of size analysis like [2, 15, 29] to verify the termination of recursive functions, and then carry out refinement typing in a second phase. However, (as confirmed by our evaluation) proving that sizes decrease often requires auxiliary invariants of the kind refinements are supposed to establish in the first place [8]. Instead, like [37], we develop a means of encoding sizes and proving termination circularly, via refinement types.

Using The Value As The Size Consider the function

```
fib  :: Nat -> Int
fib 0 = 1
fib 1 = 1
fib n = fib (n-1) + fib (n-2)
```

Recall that in our setting `Nat` is simply non-negative `Int`eger. Thus, we can naturally associate a size with its *own value*.

Termination via Environment-Weakening To verify *safety*, a standard refinement type checker (or Hoare-logic based program verifier) would check the body *assuming* an environment:

$$\texttt{n} : \texttt{Nat},\ \texttt{fib} : \texttt{Nat} \rightarrow \texttt{Int}$$

Consequently, it would verify that at the recursive call-site the argument $\texttt{n} - 2$ is a `Nat`, by checking the SMT validity of

$$(\texttt{n} \geq 0 \wedge\ \texttt{n} \neq 0 \wedge \texttt{n} \neq 1) \Rightarrow\ (\texttt{v} = \texttt{n} - 2) \Rightarrow \texttt{v} \geq 0$$

and the corresponding formula for $\mathtt{n}-1$, thereby guaranteeing that `fib` respects its signature. For *termination*, we tweak the procedure by checking the body in a *termination-weakened environment*

$$\mathtt{n} : \mathtt{Nat},\ \mathtt{fib} : \{\mathtt{n}' {:}\mathtt{Nat} \mid \mathtt{n}' < \mathtt{n}\} \rightarrow \mathtt{Int}$$

where we have weakened the type of `fib` by stipulating that it *only* be recursively called with `Nat` values $\mathtt{n}'$ that are *strictly less than* the current parameter `n`. The body still type checks as

$$(\mathtt{n} \geq 0 \wedge\ \mathtt{n} \neq 0 \wedge \mathtt{n} \neq 1)\ \Rightarrow\ (\mathtt{v} = \mathtt{n} - 2) \Rightarrow \mathtt{v} \geq 0 \wedge \mathtt{v} < \mathtt{n}$$

is a valid SMT formula. We prove (Theorem 3) that since the body typechecks under the weakened assumption for the recursive binder, the function will terminate on all `Nat`s.

Lexicographic Termination Our environment-weakening technique generalizes to include so-called lexicographically decreasing measures. For example, consider the Ackermann function.

```
ack :: Nat -> Nat -> Nat
ack m n
  | m == 0    = n + 1
  | n == 0    = ack (m-1) 1
  | otherwise = ack (m-1) (ack m (n-1))
```

We cannot prove that some argument always decreases; the function terminates because *either* the first argument strictly decreases *or* the first argument remains the same and the second argument strictly decreases, *i.e.* the pair of arguments strictly decreases according to a well-founded lexicographic ordering. To account for such functions, we generalize the notion of weakening to allow a sequence of witness arguments while requiring that (1) each *witness* argument be *non-increasing*, and (2) the *last witness* argument be *strictly decreasing* if the previous arguments were equal. These requirements can be encoded by generalizing the notion of environment-weakening: we check the body of `ack` under

$$\begin{array}{l}\mathtt{m} : \mathtt{Nat},\\ \mathtt{n} : \mathtt{Nat},\\ \mathtt{ack} : \mathtt{m}' {:}\mathtt{Nat} \rightarrow \{\mathtt{n}' {:}\mathtt{Nat} \mid \mathtt{m}' \leq \mathtt{m} \vee \mathtt{m}' = \mathtt{m} \Rightarrow \mathtt{n}' < \mathtt{n}\} \rightarrow \mathtt{Nat}\end{array}$$

Thus, while sound refinement typing requires proving termination, on the bright side, refinements make proving termination easy.

Measuring The Size of Structures Consider the function `map` defined over the standard list type.

```
map f []       = []
map f ys@(x:xs) = f x : map f xs
```

In `map`, the recursive call is made to a "smaller" input. We formalize the notion of size with *measures*.

```
measure len :: [a] -> Nat
len []     = 0
len (x:xs') = 1 + (len xs')
```

With the above definition, the **`measure`** strengthens the type of the data constructors to:

```
[]  :: {v: [a] | len v = 0}
(:) :: x:a -> xs:[a]
    -> {v:[a] | len v = 1 + len xs}
```

where `len` is simply an uninterpreted function in SMT logic [18]. We can now verify that `map` does not change the `len`gth of the list:

```
type ListEq a YS = {v:[a] | len v = len YS}

map :: (a -> b) -> ys:[a] -> ListEq b ys
```

This type is only valid if `map` provably terminates. We simultaneously verify termination and the type as before, by checking the body in the termination-weakened environment

$$\begin{array}{l}\mathtt{ys} : [\mathtt{a}]\\ \mathtt{map} : (\mathtt{a} \rightarrow \mathtt{b}) \rightarrow \{\mathtt{ys}' {:}[\mathtt{a}] \mid \mathtt{len}\ \mathtt{ys}' < \mathtt{len}\ \mathtt{ys}\} \rightarrow \mathtt{ListEq}\ \mathtt{b}\ \mathtt{ys}'\end{array}$$

To ensure that the body recursively uses `map` per its weakened specification, the recursive call `map f xs` generates the subtyping query

$$\begin{array}{l}\mathtt{ys} : \{\mathtt{len}\ \mathtt{ys} = 1 + \mathtt{len}\ \mathtt{xs}\}\\ \mathtt{xs} : \{\mathtt{len}\ \mathtt{xs} \geq 0\} \vdash \{\mathtt{ys}' = \mathtt{xs}\} \preceq \{\mathtt{len}\ \mathtt{ys}' < \mathtt{len}\ \mathtt{ys}\}\end{array}$$

Thanks to the `case` unfolding, the type of `ys` is strengthened with the measure relationship with the tail `xs` [18]. Hence, the subtyping above reduces to the valid SMT query

$$\begin{array}{l}\mathtt{len}\ \mathtt{ys} = 1 + \mathtt{len}\ \mathtt{xs}\\ \wedge\ \mathtt{len}\ \mathtt{xs} \geq 0\ \Rightarrow\ (\mathtt{ys}' = \mathtt{xs})\ \Rightarrow\ (\mathtt{len}\ \mathtt{ys}' < \mathtt{len}\ \mathtt{ys})\end{array}$$

Hence, using refinements and environment-weakening, we simultaneously verify termination and the output type for `map`.

Witnessing Termination Sometimes, the decreasing metric cannot be associated with a single parameter or lexicographically ordered sequence of parameters, but is instead an auxiliary value that is a *function* of the parameters. For example, here is the standard `merge` function from the eponymous sorting procedure:

```
merge xs@(x:xs') ys@(y:ys')
  | x < y     = x : merge xs' ys
  | otherwise = y : merge xs  ys'
```

neither parameter provably decreases in both calls, but the *sum* of the sizes of the parameters strictly decreases.

In these cases, we just explicate the metric with a *ghost* parameter that acts as a witness for termination:

```
merge d xs@(x:xs') ys@(y:ys')
  | x < y     = x : merge dx xs' ys
  | otherwise = y : merge dy xs  ys'
  where
    dx        = length xs' + length ys
    dy        = length xs  + length ys'
```

where `length :: zs:[a] -> {v:Nat | v = len zs}` returns the number of elements in the list. Now, the system verifies that `d` equals the sum of the sizes of the two lists, and that in each recursive call `dx` and `dy` are strictly smaller than `d`, thereby proving that `merge` terminates.

3. Language

Next, we present a core calculus $\lambda_{\downarrow}$ that formalizes our approach of refinement types under lazy evaluation. Instead of proving soundness directly on lazy, *call-by-name* (CBN) semantics, we prove soundness with respect to an *optimistic* (OPT) semantics where (provably) terminating terms are *eagerly* evaluated, and we separately prove an equivalence relating the CBN and OPT evaluation strategies. With this in mind, let us see the syntax (§ 3.1), the dynamic semantics (§ 3.2), and finally, the static semantics (§ 4).

3.1 Syntax

Figure 1 summarizes the syntax of expressions and types.

Constants The primitive constants include basic values like `true`, `false`, `0`, `1`, *etc.*, arithmetic and logical operators like $+, -, \leq, /$, $\wedge$, $\neg$. In addition, we include a special (untypable) `crash` constant that models errors. Primitive operations will return a `crash` when

$$
\begin{array}{rrcl}
\textbf{\textit{Expressions}} & e & ::= & x \mid c \mid \lambda x.e \mid e\ e \\
 & & \mid & \texttt{let}\ x = e\ \texttt{in}\ e \mid \mu f.\lambda x.e \\
\textbf{\textit{Basic Types}} & b & ::= & \texttt{nat} \mid \texttt{bool} \\
\textbf{\textit{Label}} & l & ::= & \uparrow \mid \downarrow \\
\textbf{\textit{Type}} & \tau & ::= & \{v{:}b^l \mid e\} \mid x{:}\tau \to \tau
\end{array}
$$

Figure 1. Syntax of $\lambda_{\downarrow}$ Terms and Types

$$
\begin{array}{rrcl}
\textbf{\textit{Contexts}} & C & ::= & \bullet \mid v\ C \mid \texttt{let}\ x = C\ \texttt{in}\ e \\
\textbf{\textit{Values}} & v & ::= & c \mid \lambda x.e \mid \mu f.\lambda x.e
\end{array}
$$

$$
\begin{array}{rcll}
C[e] & \hookrightarrow & C[e'] & \text{if } e \in \mathsf{Fin}, e \hookrightarrow e' \\
(\lambda x.e)\ e_x & \hookrightarrow & e\,[e_x/x] & \text{if } e_x \notin \mathsf{Fin} \\
(\mu f.\lambda x.e)\ e_x & \hookrightarrow & e\,[e_x/x]\,[\mu f.\lambda x.e/f] & \text{if } e_x \notin \mathsf{Fin} \\
\texttt{let}\ x = e_x\ \texttt{in}\ e & \hookrightarrow & e\,[e_x/x] & \text{if } e_x \notin \mathsf{Fin} \\
e_1\ e_2 & \hookrightarrow & e_1'\ e_2 & \text{if } e_1 \hookrightarrow e_1' \\
c\ v & \hookrightarrow & [\![c]\!](v) &
\end{array}
$$

Figure 2. Operational Semantics of $\lambda_{\downarrow}$

invoked with inputs outside their domain, *e.g.* when / is invoked with a 0 divisor, or an `assert` is called with `false`.

Expressions In addition to the primitive constants c, $\lambda_{\downarrow}$ expressions include the variables x, λ-abstractions $\lambda x.e$ and applications $e\ e$, let-binders $\texttt{let}\ x = e\ \texttt{in}\ e$, and the fix operator $\mu f.\lambda x.e$ for defining potentially diverging recursive functions.

3.2 Dynamic Semantics

We define the dynamic behavior of $\lambda_{\downarrow}$ programs in Figure 2 using a small-step contextual operational semantics.

Forcing Evaluation In our rules, contexts and values are defined as usual. In addition to the usual lazy semantics, our soundness proof requires a "helper" optimistic semantics. Thus, we parameterize the small step rules with a *force predicate* $e \in \mathsf{Fin}$ that is used to determine whether to force evaluation of the expression e or to defer evaluation until the value is *needed*.

Evaluation Rules We have structured the small-step rules so that if a function parameter or let-binder $e \in \mathsf{Fin}$ then we eagerly force evaluation of e (the first rule), and otherwise we (lazily) substitute the relevant binder with the unevaluated parameter expression (the second, third and fourth rule). The fifth rule evaluates the function in an application, as its value is always needed.

Constants The final rule, application of a constant, requires the argument be reduced to a value; in a single step the expression is reduced to the output of the primitive constant operation.

Eager and Lazy Evaluation To understand the role played by Fin, consider an expression of the form $(\lambda x.e)\ e_x$. If $e_x \in \mathsf{Fin}$ holds, then we eagerly force evaluation of e_x (via the first rule), otherwise we β-reduce by substituting e_x inside the body e (via the second rule). Note that if e_x is already a value, then trivially, only the second rule can be applied. The same idea generalizes to the *let* and *fix* cases. Thus, we get lazy (call-by-name) and eager (call-by-value) semantics by instantiating Fin appropriately

$$
\begin{array}{rcll}
e \in \mathsf{Fin} & \Leftrightarrow & \mathit{false} & \textit{Call-By-Name} \\
e \in \mathsf{Fin} & \Leftrightarrow & e \text{ not a value} & \textit{Call-By-Value}
\end{array}
$$

We write $\hookrightarrow_v$ (resp. $\hookrightarrow_n$) for the small-step relation obtained via the CBV and CBN instantiations of Fin. We write $e \hookrightarrow^j e'$ (resp. $e \hookrightarrow^* e'$) if e reduces to e' in at most j (resp. finitely many) steps.

Serious and Trivial Expressions Following Reynolds [26], we say that an expression e is *trivial* if $e \hookrightarrow_n^* v$ for some value $v \neq \texttt{crash}$. otherwise we call the expression *serious*. That is, an expression is trivial *iff* it reduces to a non-`crash` value under CBN. Note the set of trivial expressions is closed under evaluation, *i.e.* if e is trivial and $e \hookrightarrow e'$ (for any definition of Fin) then e' is also trivial.

Optimistic Evaluation We define *optimistic evaluation* as the small-step relation $\hookrightarrow_o$ obtained by defining

$$e \in \mathsf{Fin} \Leftrightarrow e \text{ is trivial} \wedge e \text{ not a value}$$

In § 4 we will prove soundness of refinement typing with respect to optimistic evaluation. To show soundness under lazy evaluation, we additionally proved [31] the following theorem that relates lazy and optimistic evaluation.

Theorem 1 (Optimistic Equivalence)**.** $e \hookrightarrow_n^* c \Leftrightarrow e \hookrightarrow_o^* c$.

4. Type System

Next, we develop the static semantics for $\lambda_{\downarrow}$ and prove soundness with respect to the OPT, and hence CBN, evaluation strategies. We develop the type system in three steps. First, we present a general type system where termination is tracked abstractly via labels and show how to achieve sound refinement typing using a generic *termination oracle* (§ 4.1). Next, we describe a concrete instantiation of the oracle using refinements (§ 4.2). This decoupling allows us to make explicit the exact requirements for soundness, and also has the pragmatic benefit of leaving open the door for employing other kinds of termination analyses.

Basic Types Figure 1 summarizes the syntax of $\lambda_{\downarrow}$ types. Basic types in $\lambda_{\downarrow}$ are natural numbers `nat` and booleans `bool`. We include only two basic types for simplicity; it is straightforward to add any type whose values can be sized, and hence ordered using a well-founded relation, as discussed in § 2.3.

Labels We use two labels to distinguish terminating and diverging terms. Intuitively, the label $\downarrow$ appears in types that describe expressions that always terminate, while the label $\uparrow$ implies that the expressions may not terminate.

Types Types in $\lambda_{\downarrow}$ include the standard basic refinement type b annotated with a label. The type $\{v{:}b^l \mid e\}$ describes expressions of type b that satisfy the refinement e. Intuitively, if the label l is $\downarrow$ then the expression definitely terminates, otherwise it may diverge. Finally, $\lambda_{\downarrow}$ types include dependent function types $x{:}\tau \to \tau$.

Safety We formalize safety properties by giving various primitive constants the appropriate refinement types. For example,

$$
\begin{array}{lcl}
\texttt{3} & : & \{v{:}\texttt{nat}^{\downarrow} \mid v = 3\} \\
(+) & : & x{:}\texttt{nat}^{\downarrow} \to y{:}\texttt{nat}^{\downarrow} \to \{v{:}\texttt{nat}^{\downarrow} \mid v = x + y\} \\
(/) & : & x{:}\texttt{nat}^{\downarrow} \to \{v{:}\texttt{nat}^{\downarrow} \mid v \neq 0\} \to \texttt{nat}^{\downarrow} \\
\texttt{error}_{\tau} & : & \{v{:}\texttt{nat}^{\downarrow} \mid \texttt{false}\} \to \tau
\end{array}
$$

We assume that for any constant c with type

$$\mathsf{Ty}(c) \doteq x{:}\{v{:}b^l \mid e\} \to \tau$$

for any value v, if the formula e is valid then the value $[\![c]\!](v)$ is not equal to `crash` and can be typed as τ. Thus, we define *safety* by requiring that a term does not evaluate to `crash`.

Serious and Trivial Types A type τ is serious if $\tau \doteq \{v{:}b^{\uparrow} \mid e\}$, otherwise, it is trivial.

Notation We write b^l to abbreviate the *unrefined* type $\{v{:}b^l \mid \texttt{true}\}$. We ensure that serious types (which are, informally speaking, assigned to potentially diverging terms) are unrefined. We write b for b^l, when label l can be either $\downarrow$ or $\uparrow$.

Well-Formedness $\boxed{\Gamma \vdash \tau}$

$$\frac{\Gamma' = \mathsf{Trivial}(\Gamma, v{:}b^{\downarrow}) \quad \Gamma' \vdash e : \texttt{bool}}{\Gamma \vdash \{v{:}b^{\downarrow} \mid e\}}\ \text{WF-}\downarrow$$

$$\frac{}{\Gamma \vdash \{v{:}b^{\uparrow} \mid \texttt{true}\}}\ \text{WF-}\uparrow \qquad \frac{\Gamma \vdash \tau \quad \Gamma, x{:}\tau \vdash \tau'}{\Gamma \vdash x{:}\tau \to \tau'}\ \text{WF-Fun}$$

Subtyping $\boxed{\Gamma \vdash \tau_1 \preceq \tau_2}$

$$\frac{\mathsf{SmtValid}(\llbracket\Gamma\rrbracket \Rightarrow \llbracket e_1\rrbracket \Rightarrow \llbracket e_2\rrbracket)}{\Gamma \vdash \{v{:}b^{\downarrow} \mid e_1\} \preceq \{v{:}b^{\downarrow} \mid e_2\}}\ \preceq\text{-}\downarrow$$

$$\frac{}{\Gamma \vdash \{v{:}b^{l} \mid e\} \preceq \{v{:}b^{\uparrow} \mid \texttt{true}\}}\ \preceq\text{-}\uparrow$$

$$\frac{\Gamma \vdash \tau_2 \preceq \tau_1 \quad \Gamma, x{:}\tau_2 \vdash \tau_1' \preceq \tau_2'}{\Gamma \vdash x{:}\tau_1 \to \tau_1' \preceq x{:}\tau_2 \to \tau_2'}\ \preceq\text{-Fun}$$

Typing $\boxed{\Gamma \vdash e : \tau}$

$$\frac{\Gamma \vdash e : \tau_1 \quad \Gamma \vdash \tau_1 \preceq \tau_2 \quad \Gamma \vdash \tau_2}{\Gamma \vdash e : \tau_2}\ \text{T-Sub}$$

$$\frac{}{\Gamma \vdash c : \mathsf{Ty}(c)}\ \text{T-Con} \qquad \frac{x{:}\{v{:}b^{\downarrow} \mid e\} \in \Gamma}{\Gamma \vdash x : \{v{:}b^{\downarrow} \mid v = x\}}\ \text{T-}\downarrow$$

$$\frac{x{:}\tau \in \Gamma \quad \tau \neq \{v{:}b^{\downarrow} \mid e\}}{\Gamma \vdash x : \tau}\ \text{T-Var}$$

$$\frac{\Gamma, x{:}\tau_x \vdash e : \tau \quad \Gamma \vdash x{:}\tau_x \to \tau}{\Gamma \vdash (\lambda x.e) : x{:}\tau_x \to \tau}\ \text{T-Fun}$$

$$\frac{\Gamma \vdash e_1 : (x{:}\tau_x \to \tau) \quad \Gamma \vdash e_2 : \tau_x}{\Gamma \vdash e_1\ e_2 : \tau\,[e_2/x]}\ \text{T-App}$$

$$\frac{\Gamma \vdash e_x : \tau_x \quad \Gamma, x{:}\tau_x \vdash e : \tau \quad \Gamma \vdash \tau}{\Gamma \vdash \texttt{let}\ x = e_x\ \texttt{in}\ e : \tau}\ \text{T-Let}$$

$$\frac{\Gamma, x{:}\tau_x, f{:}x{:}\tau_x \to \tau \vdash e : \tau \quad \Gamma \vdash x{:}\tau_x \to \tau}{\Gamma \vdash \mu f.\lambda x.e : x{:}\tau_x \to \tau}\ \text{T-Rec}$$

Figure 3. Type-checking for $\lambda_{\downarrow}$

4.1 Type-checking

Next, we present the static semantics of $\lambda_{\downarrow}$ by describing the type-checking judgments and rules. A type environment Γ is a sequence of type bindings $x{:}\tau$. We use environments to define three kinds of rules: Well-formedness, Subtyping and Typing, which are mostly standard [3, 19]. The changes are that serious types are unrefined, and binders with serious types cannot appear in refinements.

Well-formedness A judgment $\Gamma \vdash \tau$ states that the refinements in τ are boolean values in the environment Γ restricted only to the binders whose types are *trivial*. The key rule is WF-$\downarrow$, which checks that the refinement e of a trivial basic type is a boolean value under the environment Γ' which contains all of the trivial bindings in Γ extended with the binding $v : b^{\downarrow}$. To get the trivial bindings we use the function $\mathsf{Trivial}$ defined as:

$$\begin{array}{lcll} \mathsf{Trivial}(\emptyset) & \doteq & \emptyset & \\ \mathsf{Trivial}(x{:}\tau, \Gamma) & \doteq & x{:}\tau, \mathsf{Trivial}(\Gamma) & \text{if } \tau \text{ is trivial} \\ \mathsf{Trivial}(x{:}\tau, \Gamma) & \doteq & \mathsf{Trivial}(\Gamma) & \text{otherwise} \end{array}$$

The rule WF-$\uparrow$ ensures serious types are unrefined.

Subtyping A judgment $\Gamma \vdash \tau_1 \preceq \tau_2$ states that the type τ_1 is a subtype of the type τ_2 under environment Γ. That is, informally speaking, when the free variables of τ_1 and τ_2 are bound to values described by Γ, the set of values described by τ_1 is contained in the set of values described by τ_2. The interesting rules are the two rules that check subtyping of basic types. Rule $\preceq$-$\uparrow$ checks subtyping on serious basic types, and requires that the supertype is unrefined. Rule $\preceq$-$\downarrow$ checks subtyping on trivial basic types. As usual, subtyping reduces to implication checking: the *embedding* of the environment Γ strengthened with the interpretation of e_1 in the logic should imply the interpretation of e_2 in the refinement logic. The crucial difference is in the definition of embedding. Here, we keep only the trivial binders:

$$\llbracket\Gamma\rrbracket \doteq \bigwedge \{\llbracket e\,[x/v]\rrbracket \mid x{:}\{v{:}b^{l} \mid e\} \in \mathsf{Trivial}(\Gamma)\}$$

Typing A judgment $\Gamma \vdash e : \tau$ states that the expression e has the type τ under environment Γ. That is, when the free variables in e are bound to values described by Γ, the expression e will evaluate to a value described by τ. Most of the rules are standard, we discuss only the interesting ones.

A variable expression x has a type if a binding $x{:}\tau$ exists in the environment. The rule that is used for typing depends on the structure of τ. If τ is basic and trivial, the rule T-$\downarrow$ is used which, as usual, refines the basic type with the *singleton* refinement $v = x$ [23]. Otherwise, the rule T-Var is used to type x with τ. If τ is basic and serious it should be unrefined, so its type is *not* strengthened with the singleton refinement.

The dependent application rule T-App is standard: the type of the expression $e_1\ e_2$ is τ where x is replaced with the expression e_2. Note that if τ_x is serious then e_2 may diverge and so should not appear in any type. In this case, well-formedness ensures that x will not appear inside τ, and so $\tau\,[e_2/x] = \tau$ (which does not contain e_2.) Our system allows trivial arguments to be passed to a function that expects a serious input, but not the other way around.

Soundness With a Termination Oracle We prove soundness via preservation and progress theorems with respect to OPT evaluation, assuming the existence of a *termination oracle* that assigns labels to types so that serious expressions get serious types.

Hypothesis (Termination Oracle)**.** *If $\emptyset \vdash e : \tau$ and τ is a trivial type, then e is a trivial expression.*

The proof can be found in [31], and relies on two facts: (1) We define constants so that the application of a constant is defined for every value (as required for progress) but preserves typing for values that satisfy their input refinements (as required for preservation). (2) Variables with serious types do not appear in any refinements. In particular, they do not appear in any type or environment; which makes it safe to trivially substitute them with any expression whilst preserving typing. With this, we can prove that application of a serious expressions preserves typing.

We combine preservation and progress to obtain the following result. The *crash-freedom* guarantee states that if a term e is well-typed, it will not crash under *call-by-name* semantics. From preservation and progress (and the fact that the constant `crash` has no type), we can show that a well-typed term will not crash under *optimistic* evaluation ($\hookrightarrow_o$). Hence, we can use Theorem 1 to conclude that the well-typed term cannot `crash` under *lazy* evaluation ($\hookrightarrow_n$). Similarly, type-preservation is translated from $\hookrightarrow_o$ to $\hookrightarrow_n$ by observing that a term reduces to a constant under $\hookrightarrow_o$ iff it reduces to the same constant under $\hookrightarrow_n$.

Theorem 2 (Safety)**.** *Assuming the Termination Hypothesis,*

- ***Crash-Freedom:*** *If $\emptyset \vdash e{:}\tau$ then $e \not\hookrightarrow_n^* \texttt{crash}$.*
- ***Type-Preservation:*** *If $\emptyset \vdash e{:}\tau$ and $e \hookrightarrow_n^* c$ then $\emptyset \vdash c{:}\tau$.*

$$\frac{\tau \text{ is serious} \qquad \Gamma \vdash x{:}\tau_x \to \tau \qquad \Gamma, x:\tau_x, f:y{:}\tau_x \to \tau \vdash e:\tau}{\Gamma \vdash \mu f.\lambda x.e : x{:}\tau_x \to \tau}\ \text{T-REC-}\uparrow$$

$$\frac{\tau \text{ is trivial} \qquad \Gamma \vdash x{:}\tau_x \to \tau \qquad \tau_x = \{v{:}\texttt{nat}^{\downarrow} \mid e_x\} \qquad \tau_y = \{v{:}\texttt{nat}^{\downarrow} \mid e_x \wedge v < x\} \qquad \Gamma, x:\tau_x, f:y{:}\tau_y \to \tau\,[y/x] \vdash e:\tau}{\Gamma \vdash \mu f.\lambda x.e : x{:}\tau_x \to \tau}\ \text{T-REC-}\downarrow$$

Figure 4. Typing rules for Termination

4.2 Termination Analysis

Finally, we present an instantiation of the termination oracle that essentially generalizes the notion of sized types [2, 15, 29] to the refinement setting, and hence complete the description of a sound refinement type system for $\lambda_{\downarrow}$. As described in § 2, the key idea is to change the rule for typing recursive functions, so that the body of the function is checked under a termination-weakened environment where the function can only be called on *smaller* inputs.

Termination-Weakened Environments We formalize this idea by spliting the rule T-REC that types fixpoints into the two rules shown in Figure 4. The rule T-REC-↑ can only be applied when the output type τ is serious, meaning that a recursive function typed with this rule can (when invoked) diverge. In contrast, rule T-REC-↓ is used to type a recursive function that always terminates, *i.e.* whose output type is trivial. The rule requires the argument x of such a function to be a trivial (terminating) natural number. Furthermore, when typing the body e, the type of the recursive function $f : y{:}\tau_y \to \tau$ is *weakened* to enforce that the function is invoked with an argument y that is strictly smaller than x at each recursive call-site. For clarity of exposition, we require that the decreasing metric be the value of the first parameter. It is straightforward to generalize the requirement to any well-founded metric on the arguments.

We prove that our modified type system satisfies the termination oracle hypothesis. That is, in $\lambda_{\downarrow}$ trivial types are only ascribed to trivial expressions. This discharges the Termination Hypothesis needed by Theorem 2 and proves safety of $\lambda_{\downarrow}$.

Theorem 3 (Termination)**.** *If $\emptyset \vdash e{:}\tau$ and τ is trivial then e is trivial.*

The full proof is in [31], here we summarize the key parts.

Well-formed Terms We call a term e *well-formed* with respect to a type τ, written $\tau \models e$, if: (1) if $\tau \equiv \{v{:}b^{\downarrow} \mid e'\}$ then $e \hookrightarrow_n^* v$, for some value v, and, (2) if $\tau \equiv x{:}\tau_x' \to \tau'$ then $e \hookrightarrow_n^* v$, for some value v, and for all e_x such that $\emptyset \vdash e_x{:}\tau_x'$ and $\tau_x' \models e_x$, we have $\tau'\,[e_x/x] \models e\ e_x$.

Well-formed Substitutions A *substitution* θ is either empty ($\emptyset$) or of the form $\theta; [e/x]$. A substitution θ is *well-formed* with respect to an environment Γ, written $\Gamma \models \theta$, if either both are empty, or the environment and the substitution are respectively $\Gamma; x{:}\tau$ and $\theta\,[e/x]$, and (1) $\Gamma \models \theta$, (2) $\emptyset \vdash \theta e{:}\theta\tau$, and (3) $\theta\tau \models \theta e$. Now, we can connect refinement types and termination.

Lemma 1 (Termination)**.** *If $\Gamma \vdash e : \tau$ and $\Gamma \models \theta$, then $\theta\tau \models \theta e$.*

The Termination Theorem 3 is an immediate corollary of Lemma 1 where Γ is the empty environment. Since we are using refinements to prove termination, our termination proof requires soundness and vice versa. We resolve this circularity by proving Preservation, Progress, and the Termination Lemma by mutual induction, to obtain Theorem 2 *without* the Termination Hypothesis [31].

5. Evaluation

We implemented our technique by extending LIQUIDHASKELL [34]. Next, we describe the tool, the benchmarks, and a quantitative summary of our results. We then present a qualitative discussion of how LIQUIDHASKELL was used to verify safety, termination, and functional correctness properties of a large library, and discuss the strengths and limitations unearthed by the study.

Implementation LIQUIDHASKELL takes as input: (1) A Haskell *source* file, (2) Refinement type *specifications*, including refined datatype definitions, measures, predicate and type aliases, and function signatures, and (3) Predicate fragments called *qualifiers* which are used to infer refinement types using the abstract interpretation framework of Liquid typing [27]. The verifier returns as output, SAFE or UNSAFE, depending on whether the code meets the specifications or not, and, importantly for debugging the code (or specification!) the inferred types for all sub-expressions.

5.1 Benchmarks and Results

Our goal was to use LIQUIDHASKELL to verify a suite of real-world Haskell programs, to evaluate whether our approach is *efficient* enough for large programs, *expressive* enough to specify key correctness properties, and *precise* enough to verify idiomatic Haskell codes. Thus, we used these libraries as benchmarks:

- `GHC.List` and `Data.List`, which together implement many standard list operations; we verify various size related properties,
- `Data.Set.Splay`, which implements a splay-tree based functional set data type; we verify that all interface functions terminate and return well ordered trees,
- `Data.Map.Base`, which implements a functional map data type; we verify that all interface functions terminate and return binary-search ordered trees [34],
- VECTOR-ALGORITHMS, which includes a suite of "imperative" (*i.e.* monadic) array-based sorting algorithms; we verify the correctness of vector accessing, indexing, and slicing.
- BYTESTRING, a library for manipulating byte arrays, we verify termination, low-level memory safety, and high-level functional correctness properties,
- TEXT, a library for high-performance unicode text processing; we verify various pointer safety and functional correctness properties (5.2), during which we find a subtle bug.

We chose these benchmarks as they represent a wide spectrum of idiomatic Haskell codes: the first three are widely used libraries based on recursive data structures, the fourth and fifth perform subtle, low-level arithmetic manipulation of array indices and pointers, and the last is a rich, high-level library with sophisticated application-specific invariants. These last three libraries are especially representative as they pervasively intermingle high level abstractions like higher-order loops, folds, and fusion, with low-level pointer manipulations in order to deliver high-performance. They are an appealing target for LIQUIDHASKELL, as refinement types are an ideal way to statically enforce critical invariants that are outside the scope of run-time checking as even Haskell's highly expressive type system.

Results Table 1 summarizes our experiments, which covered 39 modules totaling 10204 non-comment lines of source code and 1652 lines of specifications. The results are on a machine with an Intel Xeon X5660 and 32GB of RAM (no benchmark required more than 1GB.) The upshot is that LIQUIDHASKELL is very effective on real-world code bases. The total overhead due to hints, *i.e.* the sum of **Annot**, **Qualif**, and **Wit** (**Hint** is included in **An-**

Module	LOC	Specs	Annot	Qualif	Rec	Serious	Hint	Wit	Time (s)
GHC.List	310	29 / 34	8 / 8	3 / 3	37	5	1	0	23
Data.List	504	10 / 18	7 / 7	0 / 0	50	2	5	3	38
Data.Set.Splay	149	27 / 37	5 / 5	0 / 0	17	0	4	3	25
Data.Map.Base	1396	123 / 173	12 / 12	0 / 0	88	0	9	2	219
Data.ByteString	1035	94 / 117	10 / 10	5 / 10	49	2	7	21	249
Data.ByteString.Char8	503	20 / 26	5 / 5	0 / 0	9	0	5	2	19
Data.ByteString.Fusion	447	27 / 42	0 / 0	8 / 8	23	0	0	6	55
Data.ByteString.Internal	272	44 / 101	1 / 1	10 / 10	5	0	0	0	12
Data.ByteString.Lazy	673	69 / 89	24 / 24	24 / 96	59	4	20	1	338
Data.ByteString.Lazy.Char8	406	13 / 19	4 / 4	0 / 0	7	0	4	0	20
Data.ByteString.Lazy.Internal	55	18 / 43	2 / 2	0 / 0	8	0	0	0	2
Data.ByteString.Unsafe	110	13 / 19	0 / 0	0 / 0	0	0	0	0	3
Data.Text	801	73 / 153	6 / 6	1 / 5	35	0	3	11	221
Data.Text.Array	167	35 / 87	2 / 2	7 / 9	1	0	0	1	7
Data.Text.Encoding	189	8 / 25	1 / 1	8 / 10	4	1	0	3	197
Data.Text.Foreign	84	7 / 11	0 / 0	2 / 2	7	0	0	2	6
Data.Text.Fusion	177	2 / 2	3 / 3	9 / 22	5	3	0	0	106
Data.Text.Fusion.Size	112	10 / 22	1 / 1	2 / 2	7	0	0	0	6
Data.Text.Internal	55	22 / 49	10 / 12	9 / 23	0	0	0	0	5
Data.Text.Lazy	801	74 / 150	12 / 12	1 / 5	50	0	12	1	215
Data.Text.Lazy.Builder	147	6 / 31	1 / 1	2 / 4	4	0	1	1	28
Data.Text.Lazy.Encoding	127	2 / 2	0 / 0	0 / 0	4	0	0	0	11
Data.Text.Lazy.Fusion	84	4 / 11	1 / 1	0 / 0	3	1	0	0	9
Data.Text.Lazy.Internal	66	19 / 38	5 / 5	2 / 4	5	0	0	0	3
Data.Text.Lazy.Search	104	13 / 68	5 / 5	0 / 0	6	0	5	0	161
Data.Text.Private	25	2 / 4	0 / 0	0 / 0	1	0	0	1	2
Data.Text.Search	61	5 / 15	0 / 0	2 / 2	4	0	0	3	20
Data.Text.Unsafe	74	10 / 33	5 / 5	2 / 7	0	0	0	0	6
Data.Text.UnsafeChar	51	8 / 11	0 / 0	2 / 2	0	0	0	0	4
Data.Vector.Algorithms.AmericanFlag	270	7 / 17	2 / 2	3 / 3	7	0	2	7	44
Data.Vector.Algorithms.Combinators	26	0 / 0	0 / 0	0 / 0	0	0	0	0	0
Data.Vector.Algorithms.Common	33	25 / 82	0 / 0	9 / 9	1	0	0	0	3
Data.Vector.Algorithms.Heap	170	11 / 31	1 / 1	0 / 0	4	0	1	2	53
Data.Vector.Algorithms.Insertion	51	4 / 14	0 / 0	0 / 0	2	0	0	1	4
Data.Vector.Algorithms.Intro	142	13 / 35	3 / 3	0 / 0	6	0	3	2	26
Data.Vector.Algorithms.Merge	71	0 / 0	3 / 3	1 / 1	4	0	3	4	25
Data.Vector.Algorithms.Optimal	183	3 / 12	0 / 0	0 / 0	0	0	0	0	36
Data.Vector.Algorithms.Radix	195	4 / 16	0 / 0	0 / 0	5	0	0	3	8
Data.Vector.Algorithms.Search	78	3 / 15	0 / 0	0 / 0	3	0	0	3	10
Total	10204	857 / 1652	139 / 141	112 / 237	520	18	85	83	2240

Table 1. A quantitative evaluation of our experiments. **LOC** is the number of non-comment lines of source code as reported by `sloccount`. **Specs** is the number (/ line-count) of type specifications and aliases, data declarations, and measures provided. **Annot** is the number (/ line-count) of other annotations provided, these include invariants and hints for the termination checker. **Qualif** is the number (/ line-count) of provided qualifiers. **Rec** is the number of recursive functions in the module. **Serious** is the number of functions marked as potentially non-terminating. **Hint** is the number of termination hints given to LIQUIDHASKELL, which specify *which* parameter decreases (by default: the first parameter that has a size metric). **Wit** is the number of functions that required the addition of a ghost termination witness parameter. **Time** is the time, in seconds, required to run LIQUIDHASKELL on the module.

not), is 4.5% of **LOC**. The specifications themselves are machine checkable versions of the comments placed around functions describing safe usage and behavior. Our default metric, namely the first parameter with an associated size measure, suffices to prove 67% of (recursive) functions terminating. 30% require a hint (*i.e.* the position of the decreasing argument) or a witness (3% required both), and the remaining 3% were marked as potentially diverging. Of the 18 functions marked as potentially diverging, we suspect 6 actually terminate but were unable to prove so. While there is much room for improving the running times, the tool is fast enough to be used interactively, verify a handful of API functions and associated helpers in isolation.

5.2 Case Study: Text

Next, to give a qualitative sense of the kinds of properties analyzed during the course of our evaluation, we present a brief overview of the verification of TEXT, which is the standard library used for serious unicode text processing in Haskell.

TEXT uses byte arrays and stream fusion to guarantee performance while providing a high-level API. In our evaluation of LIQUIDHASKELL on TEXT [22], we focused on two types of properties: (1) the safety of array index and write operations, and (2) the functional correctness of the top-level API. These are both made more interesting by the fact that TEXT internally encodes characters using UTF-16, in which characters are stored in either two or four bytes. TEXT is a vast library spanning 39 modules and 5700 lines of code, however we focus on the 17 modules that are relevant to the above properties. While we have verified exact functional correctness size properties for the top-level API, we focus here on the low-level functions and interaction with unicode.

Arrays and Texts A `Text` consists of an (immutable) `Array` of 16-bit words, an offset into the `Array`, and a length describing the number of `Word16`s in the `Text`. The `Array` is created and filled using a *mutable* `MArray`. All write operations in TEXT are performed on `MArray`s in the `ST` monad, but they are *frozen* into `Array`s before being used by the `Text` constructor. We write a measure denoting the size of an `MArray` and use it to type the write and freeze operations.

```
measure malen       :: MArray s -> Int
predicate EqLen A MA = alen A = malen MA
predicate Ok I A     = 0 <= I < malen A
type VO A            = {v:Int| Ok v A}

unsafeWrite  :: m:MArray s
             -> VO m -> Word16 -> ST s ()
unsafeFreeze :: m:MArray s
             -> ST s {v:Array | EqLen v m}
```

Reasoning about Unicode The function `writeChar` (abbreviating `UnsafeChar.unsafeWrite`) writes a `Char` into an `MArray`. TEXT uses UTF-16 to represent characters internally, meaning that every `Char` will be encoded using two or four bytes (one or two `Word16`s).

```
writeChar marr i c
    | n < 0x10000 = do
        unsafeWrite marr i (fromIntegral n)
        return 1
    | otherwise = do
        unsafeWrite marr i lo
        unsafeWrite marr (i+1) hi
        return 2
    where n = ord c
          m = n - 0x10000
          lo = fromIntegral
             $ (m `shiftR` 10) + 0xD800
          hi = fromIntegral
             $ (m .&. 0x3FF) + 0xDC00
```

The UTF-16 encoding complicates the specification of the function as we cannot simply require `i` to be less than the length of `marr`; if `i` were `malen marr - 1` and `c` required two `Word16`s, we would perform an out-of-bounds write. We account for this subtlety with a predicate that states there is enough `Room` to encode `c`.

```
predicate OkN I A N  = Ok (I+N-1) A
predicate Room I A C = if ord C < 0x10000
                       then OkN I A 1
                       else OkN I A 2

type OkSiz I A = {v:Nat  | OkN  I A v}
type OkChr I A = {v:Char | Room I A v}
```

`Room i marr c` says "if `c` is encoded using one `Word16`, then `i` must be less than `malen marr`, otherwise `i` must be less than `malen marr - 1`." `OkSiz I A` is an alias for a valid number of `Word16`s remaining after the index `I` of array `A`. `OkChr` specifies the `Char`s for which there is room (to write) at index `I` in array `A`. The specification for `writeChar` states that given an array `marr`, an index `i`, and a valid `Char` for which there is room at index `i`, the output is a monadic action returning the number of `Word16` occupied by the `char`.

```
writeChar :: marr:MArray s
          -> i:Nat
          -> OkChr i marr
          -> ST s (OkSiz i marr)
```

Bug Thus, clients of `writeChar` should only call it with suitable indices and characters. Using LIQUIDHASKELL we found an error in one client, `mapAccumL`, which combines a map and a fold over a `Stream`, and stores the result of the map in a `Text`. Consider the inner loop of `mapAccumL`.

```
outer arr top = loop
 where
  loop !z !s !i =
    case next0 s of
      Done          -> return (arr, (z,i))
      Skip s'       -> loop z s' i
      Yield x s'
        | j >= top  -> do
          let top' = (top + 1) `shiftL` 1
          arr' <- new top'
          copyM arr' 0 arr 0 top
          outer arr' top' z s i
        | otherwise -> do
          let (z',c) = f z x
          d <- writeChar arr i c
          loop z' s' (i+d)
      where j | ord x < 0x10000 = i
              | otherwise       = i + 1
```

Let's focus on the `Yield x s'` case. We first compute the maximum index `j` to which we will write and determine the safety of a write. If it is safe to write to `j` we call the provided function `f` on the accumulator `z` and the character `x`, and write the *resulting* character `c` into the array. However, we know nothing about `c`, in particular, whether `c` will be stored as one or two `Word16`s! Thus, LIQUIDHASKELL flags the call to `writeChar` as *unsafe*. The error can be fixed by lifting `f z x` into the `where` clause and defining the write index `j` by comparing `ord c` (not `ord x`). LIQUIDHASKELL (and the authors) readily accepted our fix.

5.3 Code Changes

Our case studies also highlighted some limitations of LIQUIDHASKELL that we will address in future work. In most cases, we could alter the code slightly to facilitate verification. We briefly summarize the important categories here; refer to [31] for details.

Ghost parameters are sometimes needed in order to materialize values that are not needed for the computation, but are necessary to prove the specification; proving termination may require a decreasing value that is a function of several parameters. In future work it will be interesting to explore the use of advanced techniques for synthesizing ranking witnesses [8] to eliminate such parameters.

Lazy binders sometimes get in the way of verification. A common pattern in Haskell code is to define *all* local variables in a single `where` clause and use them only in a subset of all branches. LIQUIDHASKELL flags a few such definitions as *unsafe*, not realizing that the values will only be demanded in a specific branch. Currently, we manually transform the code by pushing binders inwards to the usage site. This transformation could be easily automated.

Assumes which can be thought of as "hybrid" run-time checks, had to be placed in a couple of cases where the verifier loses information. One source is the introduction of assumptions about mathematical operators that are currently conservatively modeled in the refinement logic (*e.g.* that multiplication is commutative and associative). These may be removed by using more advanced nonlinear arithmetic decision procedures.

6. Related Work

Next we situate our work with closely related lines of research.

Dependent Types are the basis of many verifiers, or more generally, proof assistants. In this setting arbitrary terms may appear inside types, so to prevent logical inconsistencies, and enable the checking of type equivalence, all terms must terminate. "Full" dependently typed systems like Coq [4], Agda [21], and Idris [6] typically use various *structural* checks where recursion is allowed on sub-terms of ADTs to ensure that *all* terms terminate. One can *fake* "lightweight" dependent types in Haskell [7, 10, 24]. In this style, the invariants are expressed in a restricted [16] total index language and relationships (*e.g.* $x < y$ and $y < z$) are combined (*e.g.* $x < z$) by explicitly constructing a term denoting the consequent from terms denoting the antecedents. On the plus side this "constructive" approach ensures soundness. It is impossible to witness inconsistencies, as doing so triggers diverging computations. However, it is unclear how easy it is to use restricted indices with explicitly constructed relations to verify the complex properties needed for large libraries.

Refinement Types are a form of dependent types where invariants are encoded via a combination of types and SMT-decidable logical refinement predicates [28, 38]. Refinement types offer a highly automated means of verification and have been applied to check a variety of program properties, including functional correctness of data structures [9, 18], security protocols [3, 5] and compilers [13]. The language of refinements is restricted to ensure consistency, however, program variables (binders), or their singleton representatives [38], *are* crucially allowed in the refinements. As discussed in § 2 this is sound under *call-by-value* evaluation, but under Haskell's semantics any innocent binder can be potentially diverging, causing unsoundness. Finally, our implementation is based on LIQUIDHASKELL [34], which was unsound as it assumed CBV evaluation.

Size-based Termination Analyses have been used to verify termination of recursive functions, either using the "size-change principle" [2, 17], or via the type system [15, 29] by annotating types with size indices and verifying that the arguments of recursive calls have smaller indices. In work closely related to ours, Xi [37] encoded sizes via refinement types to prove totality of programs. What differentiates the above work from ours is that we do not aim to prove that all expressions converge; on the contrary, under a lazy setting diverging expressions are welcome. We use size analysis to track diverging terms in order to exclude them from the logic.

Static Checkers like ESCJava [12] are a classical way of verifying correctness through assertions and pre- and post-conditions. One can view Refinement Types as a type-based generalization of this approach. Classical contract checkers check "partial" (as opposed to "total") correctness (*i.e.* safety) for *eager*, typically first-order, languages and need not worry about termination. We have shown that in the lazy setting, even "partial" correctness requires proving "total" correctness! [39] describes a static contract checker for Haskell that uses symbolic execution. The (checker's) termination requires that recursive procedures only be unrolled up to some fixed depth. While this approach removes inconsistencies, it yields weaker, "bounded" soundness guarantees. Zeno [30] is another automatic prover for Haskell which proves properties by unrolling recursive definitions, rewriting, and goal-splitting, using sophisticated proof-search techniques to ensure convergence. As it is based on rewriting, "Zeno might loop forever" when faced with non-terminating functions, but will not conclude erroneous facts. Finally, [35] describes a novel contract checking technique that encodes Haskell programs into first-order logic. Intriguingly, the paper shows how the encoding, which models the denotational semantics of the code, is *simplified* by lazy evaluation.

Unlike the previous contract checkers for Haskell, our type-based approach does not rely on heuristics for unrolling recursive procedures, and instead uses SMT and abstract interpretation [27] to infer signatures, which we conjecture makes LIQUIDHASKELL more predictable. Of course, this requires LIQUIDHASKELL be provided logical qualifiers (predicate fragments) which form the basis of the analysis' abstract domain. In our experience, however, this is not an onerous burden as most qualifiers can be harvested from API specifications, and the overall workflow is predictable enough to enable the verification of large, real-world code bases.

References

[1] T. Ball, R. Majumdar, T. Millstein, and S. K. Rajamani. Automatic predicate abstraction of C programs. In *PLDI*, 2001.

[2] G. Barthe, M. J. Frade, E. Giménez, L. Pinto, and T. Uustalu. Type-based termination of recursive definitions. *Mathematical Structures in Computer Science*, 2004.

[3] J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffeis. Refinement types for secure implementations. *ACM TOPLAS*, 2011.

[4] Y. Bertot and P. Castéran. *Coq'Art: The Calculus of Inductive Constructions*. Springer Verlag, 2004.

[5] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P.-Y Strub. Implementing tls with verified cryptographic security. In *IEEE S & P*, 2013.

[6] Edwin Brady. Idris: general purpose programming with dependent types. In *PLPV*, 2013.

[7] M. T. Chakravarty, G. Keller, and S. L. Peyton-Jones. Associated type synonyms. In *ICFP*, 2005.

[8] B. Cook, A. Podelski, and A. Rybalchenko. Termination proofs for systems code. In *PLDI*, 2006.

[9] J. Dunfield. Refined typechecking with Stardust. In *PLPV*, 2007.

[10] R. A. Eisenberg and S. Weirich. Dependently typed programming with singletons. In *Haskell Symposium*, 2012.

[11] R. Ennals and S. Peyton-Jones. Optimistic evaluation: An adaptive evaluation strategy for non-strict programs. In *ICFP*, 2003.

[12] C. Flanagan, K.R.M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. In *PLDI*, 2002.

[13] C. Fournet, N. Swamy, J. Chen, P-É. Dagand, P-Y. Strub, and B. Livshits. Fully abstract compilation to javascript. In *POPL*, 2013.

[14] M. Greenberg, B. C. Pierce, and S. Weirich. Contracts made manifest. *JFP*, 2012.

[15] J. Hughes, L. Pareto, and A. Sabry. Proving the correctness of reactive systems using sized types. In *POPL*, 1996.

[16] L. Jia, J. Zhao, V. Sjöberg, and S. Weirich. Dependent types and program equivalence. In *POPL*, 2010.

[17] Neil D. Jones and Nina Bohr. Termination analysis of the untyped lamba-calculus. In *RTA*, 2004.

[18] M. Kawaguchi, P. Rondon, and R. Jhala. Type-based data structure verification. In *PLDI*, 2009.

[19] K.W. Knowles and C. Flanagan. Hybrid type checking. *ACM TOPLAS*, 2010.

[20] A. Mycroft. The theory and practice of transforming call-by-need into call-by-value. In *ESOP*, 1980.

[21] U. Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Chalmers, 2007.

[22] B. O'Sullivan and T. Harper. text-0.11.2.3: An efficient packed unicode text type. http://hackage.haskell.org/package/text-0.11.2.3.

[23] X. Ou, G. Tan, Y. Mandelbaum, and D. Walker. Dynamic typing with dependent types. In *IFIP TCS*, 2004.

[24] S. L. Peyton-Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple unification-based type inference for GADTs. In *ICFP*, 2006.

[25] Gordon Plotkin. Call-by-name, call-by-value and the lambda calculus. *Theoretical Computer Science*, 1975.

[26] John C. Reynolds. Definitional interpreters for higher-order programming languages. In *25th ACM National Conference*, 1972.

[27] P. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In *PLDI*, 2008.

[28] J. Rushby, S. Owre, and N. Shankar. Subtypes for specifications: Predicate subtyping in pvs. *IEEE TSE*, 1998.

[29] D. Sereni and N.D. Jones. Termination analysis of higher-order functional programs. In *APLAS*, 2005.

[30] W. Sonnex, S. Drossopoulou, and S. Eisenbach. Zeno: An automated prover for properties of recursive data structures. In *TACAS*, 2012.

[31] From Safety to Termination and back: SMT-Based verification for Lazy languages. Supplementary Materials.

[32] A. M. Turing. On computable numbers, with an application to the eintscheidungsproblem. In *LMS*, 1936.

[33] H. Unno, T. Terauchi, and N. Kobayashi. Automating relatively complete verification of higher-order functional programs. In *POPL*, 2013.

[34] N. Vazou, P. Rondon, and R. Jhala. Abstract refinement types. In *ESOP*, 2013.

[35] D. Vytiniotis, S.L. Peyton-Jones, K. Claessen, and D. Rosén. Halo: haskell to logic through denotational semantics. In *POPL*, 2013.

[36] P. Wadler. Call-by-value is dual to call-by-name. In *ICFP*, 2003.

[37] H. Xi. Dependent types for program termination verification. In *LICS*, 2001.

[38] H. Xi and F. Pfenning. Eliminating array bound checking through dependent types. In *PLDI*, 1998.

[39] Dana N. Xu, Simon L. Peyton-Jones, and Koen Claessen. Static contract checking for haskell. In *POPL*, 2009.